Auction Analysis by Normal Form Game Approximation

Michael Kaisers
Eindhoven University of Technology
Industrial Design Department
P.O. Box 513
5600 MB Eindhoven, The Netherlands
michaelkaisers@gmail.com

Karl Tuyls
Eindhoven University of Technology
Industrial Design Department
P.O. Box 513
5600 MB Eindhoven, The Netherlands
k.p.tuyls@tue.nl

Frank Thuijsman
Maastricht University
MICC / Mathematics
P.O. Box 616
6200 MD Maastricht, The Netherlands
frank@micc.unimaas.nl

Simon Parsons
Brooklyn College
Department of Computer and Information Science
Brooklyn College, City University of New York
Brooklyn NY 11210 USA
parsons@sci.brooklyn.cuny.edu

Abstract

Auctions are pervasive in today's society and provide a variety of markets, ranging from consumer-to-consumer online auctions to government-to-business auctions for telecommunications spectrum licenses. This article enables a strategic choice between a set of available trading strategies by introducing a methodology to approximate heuristic payoff tables by normal form games. An example from the auction domain is transformed by this means and an evolutionary game theory analysis is applied subsequently. The information loss in the normal form approximation is shown to be reasonably small such that the concise normal form representation can be leveraged in order to make strategic decisions in auctions. In particular, a mix of trading strategies that guarantees a certain profit against any population of traders is computed and further applications are indicated.

Keywords: *Evolutionary game theory, Auction theory, Multi-agent games*

1. Introduction

Auctions are deployed in a variety of real markets to foster highly efficient trading. They range from consumer-to-consumer markets like eBay and business-to-business stock exchanges to government-to-business auctions for mineral rights or government licenses for the telecommunications spectrum [6, 8]. Furthermore, auction mechanisms have been transferred successfully to solve other resource allocation problems, e.g. in the domain of efficient internet traffic routing [11]. This motivates researching how to run auctions and how to extract profit by trading within them as auctions are pervasive in today's society.

The traders that participate in an auction agree to subject to a set of market rules in order to exchange goods for money. Within the scope of this article only commodity markets are considered, i.e. a single type of an abstract good is traded. Each trader is assumed to have a *private valuation* of the good which is only known to himself. Buyers and sellers place offers to indicate their intention to trade at a certain price. The here considered *clearing house auction* proceeds in rounds and polls offers from each trader each round. When all offers are collected, an equilibrium price is established based on the available offers such that demand meets supply at this price. It is commonly set to the average of the two offers that define the range of possible equilibrium prices, i.e. the lowest bid and the highest ask that can be matched in the equilibrium. Each buyer with an offer above that price is matched with a seller having an offer below that price. The *profit* of a transaction can be computed as the difference between the transaction price and the private value, assuming that buyers will not buy above their private value and sellers will not sell below their private value.

A multitude of trading strategies has been devised to derive the next offer, possibly exploiting the knowledge about offers and transactions that were observed in previous rounds. The most trivial one is *Truth Telling* (TT) which just reveals the private value by placing offers exactly at that value. Despite its simplicity, it may be optimal in some situations [13]. The experiment of this article considers three more sophisticated trading strategies. Roth and Erev devised a reinforcement learning model of human trading behavior in [2] which is modified to perform in a clearing house auction as *Modified Roth-Erev* (MRE) [7]. MRE is evaluated in competition to *Gjerstad and Dickhaut* (GD) and *Zero Intelligence Plus* (ZIP). GD maximizes the expected profit by computing the profit and probability of leading to a transaction for a set of relevant prices [3]. ZIP places stochastic bids within a certain profit margin, which is lowered when a more competitive offer was rejected and increased when a less competitive offer was accepted [1].

Given a set of available trading strategies, it is of high interest which strategy is *best* in the sense that it yields the highest expected payoff. However, this question cannot be answered in general as the performance of a trading strategy is highly dependent on the competition it faces [12]. Let us assume an auction where traders only choose between the trading strategies described above. The profit of each trader is dependent on the overall mix of strategies and traders may choose to change their strategy in the course of time, e.g. applying a reinforcement learning algorithm to improve their expected payoff. This adaptation can be modeled by the *replicator dynamics* from evolutionary game theory which are formally connected to reinforcement learning [14].

A *heuristic payoff table* is proposed in [15] and adopted by several authors to capture the average profit of each type of trading strategy for all possible mixtures of strategies in a finite population [5, 10]. For the domain of auctions, the required profits can only be computed in simulation where the private valuation of each trader is known. This table is a first step towards revealing the dynamics of adopted trading strategies in auctions and can for example be used to analyze which trading strategy yields the highest potential for improvements [10]. Although the heuristic payoff table provides the basis for analyzing the dynamics in auctions, it is unintuitive and lacks information about the payoffs for strategies that are not yet present in a population. However, exactly these payoffs would provide information about whether it is profitable or not to be the first one to adopt this strategy. The normal form game on the other hand enables an individual trader to calculate his expected profit for each of his possible choices against any mix of strategies he faces. It is more intuitive and allows inspecting the strategic situation with means from game theory, e.g. allowing to compute optimal strategies, best replies and Nash equilibria more easily.

This suggests the question whether a heuristic payoff table can be approximated by a normal form game in order to open up these opportunities. Answering this question affirmatively, this article demonstrates how an approximation can be found using a linear least squares algorithm or linear programming. The methodology is illustrated by approximating a heuristic payoff table from the auction domain and the results presented below show a reasonably small error such that the approximation can be leveraged for strategic considerations and an intuitive grasp of the game in auctions.

The remainder of this article is structured as follows: Section 2 introduces the game theoretical background that is required by the methodology presented in Section 3. Subsequently, Section 4 illustrates the method by applying it to an example from the auction domain and presents the resulting performance. These results are discussed in Section 5 and the paper is concluded with future directions in Section 6.

2. Game theoretical background

Classical game theory is the mathematical study of strategic conflicts of rational agents. Each individual i chooses a pure strategy s_i from the set of available strategies S_i due to some strategy π_i and has a preference relation over all possible outcomes. The players are assumed to choose their actions simultaneously and independently. This implies that the preference relation can be captured by a numerical *payoff function* τ_i which is public knowledge and assigns a value of desirability to each possible joint strategy $s = (s_1, \ldots, s_n)$, where n is the total number of agents.

$$\tau_i : S_1 \times \ldots \times S_n \to \mathbb{R}$$

In the context of auctions, each pure strategy corresponds to a trading strategy and the preference relation is proportional to the profit that an agent can make given the set of opponents' trading strategies. This section introduces two different means to capture payoff functions for multi-agent games and auctions in particular: The normal form game and the heuristic payoff table. Furthermore, advantages and disadvantages of both representations are discussed. Subsequently, the concepts of replicator dynamics and basins of attraction are presented.

2.1. Normal form games

A normal form game commonly describes the payoff to each agent in matrix notation. The matrix given in Figure 1 describes a symmetric two-player normal form game where both players may choose between the three strategies Rock, Paper and Scissors. The first player may choose a row r, the second player chooses a column c and the joint choice (r, c) determines the payoff which the matrix gives for the first player. However, the payoff matrix for the second player

	$Rock$	$Paper$	$Scissors$
$Rock$	0	-1	1
$Paper$	1	0	-1
$Scissors$	-1	1	0

Figure 1. Payoffs for the row player in the symmetric two-player normal form game 'Rock-Paper-Scissors'.

equals the transposed of the first player's payoffs in symmetric games. Hence, it can be derived from the same table by consulting the entry (c, r). Both players seek to maximize their expected payoff and optimal mixed strategies can be derived such that a highest least profit v against any opponent is guaranteed by randomizing over the pure strategies. This profit v is also called the value of a game.

Within a normal form game, the player optimizes his expected payoff against an opponent that plays according to a certain probability distribution. Similarly, he faces a field of traders that are distributed over the strategies in reality. It does not actually matter which opponent plays which strategy but rather how many opponents deploy which strategy. Assuming he encounters a random individual from the population, his opponent's strategy will be drawn from the distribution of strategies in the population. The opponent in normal form games therefore resembles the population in which the agent is situated in reality.

2.2. Heuristic payoff tables

A heuristic payoff table may also be used to capture the payoffs of a game. However, it requires a finite population of traders such that all possible combinations of strategies can be evaluated. If each agent $i \in \{1, 2, \ldots, n\}$ has to choose a pure strategy $s_i \in \{1, 2, \ldots, k\}$, this leads to a joint strategy $(s_1, \ldots, s_n)$. However, for an individual trader it is only important to know how many of his opponents are playing each of the different strategies. So, given $(s_1, \ldots, s_n)$ the individual trader could derive that there are N_1 agents playing strategy 1, N_2 agents playing strategy 2, etc.. This would yield a *discrete profile* $N = (N_1, \ldots, N_k)$ telling exactly how many agents play each strategy. The average profit for playing a strategy can then be denoted by a payoff vector $U(N) = (U_1(N), \ldots, U_k(N))$ indicating that strategy $s \in \{1, 2, \ldots, k\}$ would yield an average payoff of $U_s(N)$ for the discrete profile N. The distribution of n agents on k pure strategies is a combination with repetition, hence the number of rows of a heuristic payoff table is given by:

$$\binom{n+k-1}{k}$$

The payoffs of these discrete profiles can be measured in many practical domains, including poker and auctions. However, measurements do not allow to capture the payoff to strategies that are not present, i.e. whenever $N_s = 0$ then $U_s(N)$ is unknown for that discrete profile. Table 1 shows a heuristic payoff table obtained from an auction simulation, indicating unknown payoffs with a dash.

The heuristic payoff table is an approximation of a symmetric game. The full payoff function for this game would map each joint strategy to a payoff for each agent and require k^n entries. In contrast to this, the heuristic payoff table maps discrete strategy profiles to payoffs for each strategy. The normal form game maps probability distributions over the strategies to payoffs for each strategy and only requires k^2 entries. The example presented below features 6 agents and 3 strategies and would require $3^6 = 729$ entries for the full representation, $\binom{6+3-1}{3} = 28$ rows or $28 \cdot 3 = 84$ payoff entries in the heuristic payoff table and only $3^2 = 9$ entries for the normal form game approximation which reduces the multi-player game to a two-player game.

Table 1. The heuristic payoff table of a clearing house auction with 6 agents and the three strategies ZIP, MRE and GD. The first three columns give the discrete profiles N over the trading strategies and the last three columns give the corresponding payoff vectors $U(N)$.

N_{ZIP}	N_{MRE}	N_{GD}	U_{ZIP}	U_{MRE}	U_{GD}
6	0	0	99	-	-
5	1	0	97	100	-
5	0	1	89	-	69
4	2	0	96	94	-
4	1	1	90	88	65
4	0	2	85	-	69
3	3	0	97	92	-
3	2	1	87	90	64
3	1	2	85	80	73
3	0	3	76	-	73
2	4	0	97	96	-
2	3	1	91	91	66
2	2	2	84	83	67
2	1	3	78	70	76
2	0	4	62	-	80
1	5	0	97	97	-
1	4	1	93	89	62
1	3	2	86	84	69
1	2	3	73	71	75
1	1	4	73	57	77
1	0	5	56	-	80
0	6	0	-	94	-
0	5	1	-	91	62
0	4	2	-	84	67
0	3	3	-	75	71
0	2	4	-	65	76
0	1	5	-	43	79
0	0	6	-	-	79

The heuristic payoff table is non-intuitive and needs a thorough interpretation. One opportunity to make it more accessible is the approximation of the heuristic payoff table by a normal form game which allows to draw on the classical means from game theory to derive strategic choices.

While a normal form game yields the more detailed pairwise comparison of strategies and is more accessible for analysis, it also imposes linearity over the different payoffs to profiles. The heuristic payoff table on the other hand only captures the average payoffs but these may be very non-linear over the different discrete profiles.

2.3. Replicator dynamics

Replicator dynamics describe the game dynamics from an evolutionary perspective. Evolutionary game theory assumes an infinitely large population of individuals that choose their pure strategy according to some probability distribution. It assumes this population to evolve such that successful strategies with higher payoffs grow while less successful ones decay which allows to analyze the asymptotic behavior of this population.

Evolutionary game theory takes a rather descriptive perspective replacing hyper-rationality from classical game theory by the concept of natural selection from biology. The evolutionary pressure by natural selection can be modeled by the replicator equations. This article assumes one population playing the symmetric game and hence uses the single-population replicator dynamics that define the growth of a strategy proportional to the fraction of the population that already uses this strategy and the difference between the payoff to this strategy and the average payoff.

The game dynamics for the normal form game with payoff matrix A can be calculated for strategy i given that the opponent mixes over the pure strategies according to the probability vector x:

$$\dot{x_i} = x_i \cdot [(Ax)_i - xAx]$$

It is also possible to construct the replicator dynamics from the heuristic payoff table. Assuming each agent i independently chooses his pure strategy $s_i \in \{1, \ldots, k\}$ according to some probability distribution $p = (p_1, \ldots, p_k)$, the probability of each joint strategy is $\prod_i^n p_{s_i}$. The probability of a discrete profile can be computed as the product of the number of joint strategies that lead to this discrete profile and the probability of these joint strategies. It is a multinomial for which $Pr(N|p)$ is the probability of the discrete profile N given the mixed strategy p.

$$Pr(N|p) = \binom{n}{N_1, \ldots, N_k} \cdot p_1^{N_1} \cdot \ldots \cdot p_k^{N_k}$$

The payoff for each strategy can then be computed as the weighted average over the payoffs received in all profiles. However, a correction term is required if the payoffs for non-occurring strategies is unknown in the heuristic payoff table.

$$U_{average,i} = \frac{\sum_N Pr(N|p) \cdot U_i(N)}{1 - Pr(unknown|i)}$$

Now, the game dynamics can be computed from the heuristic payoff table:

$$\dot{x_i} = x_i \cdot [U_{average,i} - x(U_{average,i})^T]$$

The resulting dynamics can be visualized in a force field plot as in Figure 3 where the arrows indicate the direction and strength of change.

The replicator dynamics give rise to a dynamical system which may feature repellers and attractors of which the latter are of particular importance to the analysis of asymptotic behavior. Each attractor consumes a certain amount of the strategy space that eventually converges to it — this space is also called the basin of attraction [4]. Assuming that an evolutionary process may start uniformly at any point in the strategy space, the size of the basin of attraction may be used to estimate the practical importance of an attractor. This can be achieved by uniform sampling of the strategy space and analysis of trajectory convergences or by inspection of the directional field plots given in Figure 3. A good approximation of the game dynamics should have minimal impact on the basins of attraction.

3. Methodology

The conversion of a normal form game to a heuristic payoff table is a quite trivial step which suggests that the inverse is also possible. However, the inverse transformation is over constrained and a heuristic payoff table can only be approximated by a normal form game. This section presents the newly proposed method for finding a suitable normal form game approximation.

3.1. From normal form games to heuristic payoff tables

The heuristic payoff table lists all possible discrete profiles with the average payoff of playing against a finite population that mixes accordingly. The payoff vector against the mixed strategy p can be computed from the game matrix M as Mp. Let D be the matrix where each row corresponds to a discrete profile N of n agents and let $P = \frac{1}{n} \cdot D$ map the discrete profiles to probabilities. The matrix U that yields the corresponding payoff vectors $U(N)$ as rows can then be computed as the product of P and M.

$$U = P \cdot M^T \tag{1}$$

The heuristic payoff table $H = (D, U)$ is the composition of the discrete profiles and the corresponding payoffs.

3.2. From heuristic payoff tables to normal form games

This section reverses the step of the previous section and shows the transition from a heuristic payoff table to a normal form game approximation. However, equation (1) cannot simply be solved for M as the values in the heuristic payoff table may be noise-prone due to stochasticity in the experiments and may also feature non-linear dynamics which leads to an over-constrained system of equations. Therefore, it needs to be approximated, e.g. by minimizing the mean squared error or the maximal absolute deviation.

Minimizing mean squared error

A normal form game M that approximates the heuristic payoff table $H = (D, U)$ can be determined incrementally for each row M_i by finding a least mean squared error fit between the i'th column of U, denoted as U_i, and the reconstructed payoff vector $\tilde{U}_i = P \cdot M_i^T$ from the normal form game, where $P = \frac{1}{n} \cdot D$ as above, by solving the minimization problem:

$$\min_{M_i} \; ||U_i - \tilde{U}_i||^2$$

A standard linear least square fitting algorithm can be used to solve this system for each row and compose the normal form game matrix.

Minimizing maximal absolute deviation

Linear programming optimizes a linear goal function subject to a system of linear inequalities. Using the same definitions of the profile matrix D, the probability matrix P, the game M and the payoff matrix U as above, the following program can be formulated.

$$\begin{aligned} \text{minimize} \;\; & \epsilon \\ \text{variables} \;\; & \epsilon, M_{ij}, \text{for } i, j \in \{1, \ldots, k\} \\ \text{subject to} \;\; & P \cdot M^T \leq U + \epsilon \\ & P \cdot M^T \geq U - \epsilon \end{aligned}$$

However, this program needs to be transformed to standard notation in order to apply common algorithms from linear programming. For sake of convenience, each row M_i is determined separately. Let $c = (1, 0, \ldots, 0)$ and $x = (\epsilon, M_i)$ such that the goal function minimizes epsilon. Furthermore, let $A = \begin{pmatrix} -1 & \\ & P \\ \vdots & -P \\ -1 & \end{pmatrix}$ and $b = \begin{pmatrix} U_i \\ -U_i \end{pmatrix}$ where U_i is the i'th column of the payoff matrix. Then, this linear program can be solved in standard notation:

$$\min_x \; c \cdot x^T \quad \text{subject to} \;\; A \cdot x^T \leq b, x \geq 0$$

In order to approximate the heuristic payoff table, we need to solve k linear programs to compute the complete normal form matrix.

4. Experiments

This section presents the experimental setup and results of measuring the information loss in the normal form game approximation of a heuristic payoff table from the auction domain. The heuristic payoff table given in Table 1 is obtained by simulating auctions with the *Java Auction Simulator API* (JASA) [9]. This empirical platform contains the trading strategies ZIP, MRE and GD which were setup with the following parameters: ZIP uses a learning rate of 0.3, a momentum of 0.05 and a JASA specific scaling of 0.2, MRE chooses between 40 discrete prices using a recency parameter of 0.1, an exploration of 0.2 and scaling of 9 and GD evaluates prices up to 360.

The heuristic payoff table is obtained from an average of 2000 iterations of clearing house auctions. On the start of each auction, all traders are initialized without knowledge of previous auctions and with a private value drawn from the same distribution as in [15], i.e. an integer lower bound b is drawn uniformly from $[61, 160]$ and the upper bound from $[b + 60, b + 209]$ for each buyer. The sellers' private values are initialized similarly. These private values then remain fixed over the course of the auction which runs 300 rounds on each of 5 trading days where each trader is entitled to trade one item per day.

The heuristic payoff table is approximated as described in Section 3.2, which leads to the normal form game representations given in Figure 2. The replicator dynamics are derived from the heuristic payoff table and the normal form game representations and compared in Figure 3. There is a clear qualitative correspondence of the dynamics that arise from the three models. Differences are very small and hard to identify from the force field plots. Therefore, directional field plots are given as well, which allows to find the attractors and basins of attraction by inspection.

A mixed attractor can be found at $(0.81, 0.18, 0.0)$ for the heuristic payoff table, at $(1, 0, 0)$ in least mean squared error fitting and at $(0.72, 0.27, 0.0)$ in minimized maximal absolute deviation. However, there are the same basins of attraction. The pure attractor at $(0, 0, 1)$ is present in all dynamics and is estimated to consume 26.0% of the strategy space in the heuristic payoff table in comparison to 29.0% and 26.9% in the approximations, based on the analysis of convergence of 1000 trajectories with uniformly sampled starting points. In the context of evolutionary game theory, evolutionary stable strategies provide a concept to find stable solutions in symmetric normal form games. The attractors are evolutionary stable in the normal form game approximations and predict the attractors that are observed in the auction game dynamics with a small error. The maximal absolute deviation is 9.76% and 6.64% while the square root of the mean squared error is 3.03% and 3.42% respectively.

Least mean squared error

	ZIP	MRE	GD
ZIP	97.4	98.8	52.3
MRE	96.8	98.6	42.6
GD	64.8	59.1	83.4

Least maximal absolute deviation

	ZIP	MRE	GD
ZIP	93.8	102.7	52.9
MRE	94.9	100.0	38.3
GD	66.2	60.5	81.8

Figure 2. The symmetric two-player normal form game approximation of the heuristic payoff table for a clearing house auction with the three strategies ZIP, MRE and GD as obtained by least mean squared error fitting (left) and minimizing maximal absolute deviation (right).

5. Discussion

The results show that heuristic payoff tables in the domain of auctions may be approximated by normal form games with a reasonably small error. However, this case study is rather a proof of concept and still needs to be verified by further theoretical analysis. Therefore, this section starts with a discussion of the limitations of this approach. Eventually, it is illustrated how the newly gained insights can be leveraged for strategic choice.

5.1. Linearity

The proposed approach is general in the number of actions and can be transfered to higher dimensions. However, the approximation of heuristic payoff tables by normal form games imposes a linear model on the data. This may be an oversimplification for complex dynamics which may arise from intricate interactions. Consequently, the precision of the approximation is likely to deteriorate when the number of trading strategies to choose from is increased.

5.2. Strategic choice

Consider the normal form representation of the auction game obtained from minimizing the maximal absolute deviation as given in Figure 2. It is possible to derive an *optimal strategy* that gives a lower bound on the profit that can be guaranteed even if nothing is known about the opponents. This profit is also known as the matrix game value, equals 73.1 for this example and can be guaranteed by the optimal trading strategy $\pi^* = (0.3, 0, 0.7)$. This means that a trader who plays ZIP with probability 0.3 and GD with probability 0.7 will get an expected payoff of at least 73.1 against any combination of GD, ZIP an MRE. For any other probability distribution than π^* he may encounter an opponent that gives him a lower expected payoff. Calculation of the matrix game value and the optimal strategy can be achieved by linear programming since the agent simply wants to choose his probabilities in such a way that the minimal payoff over all columns is maximized.

If an agent would know the actual current mix of trading strategies in the population he faces, he could even make more than with the optimal strategy because the optimal strategy is generally not a best reply against a specific population distribution. It is rather like a best worst case analysis.

6. Conclusions

This article has modeled trading in auctions by considering a population of traders that repeatedly participate in an auction. A set of trading strategies is made available to the agents who make their choice according to the relative profit of these strategies.

The contributions can be summarized as follows: A methodology to approximate heuristic payoff tables by normal form games has been introduced. This smaller game representation is more intuitive and computationally less expensive to analyze and fills in a gap of missing payoffs in the blind spots of the heuristic payoff table. In fact, the normal form game can even be constructed from partial heuristic payoff tables, e.g. when a number of profiles could not be observed. Rather than merely participating myopically, a rational agent can now inspect the game strategically and means and reasoning from game theory can be applied, e.g. to analyze asymptotic properties of the auction.

Future work will aim to argue for the described approach on a more theoretical level and look for structure in the deviation from the linear model, in particular where and why qualitative changes occur. Further interesting opportunities include extending the game theoretic analysis, e.g. investigating symmetric equilibria under different replicator dynamics that may account for exploration and deliver a more realistic model of human behavior. Furthermore, this approach needs to be tested on other auctions and domains, possibly applying it to higher dimensions as it is general in the number of strategies.

Acknowledgments

The authors wish to express their gratitude to Jinzhong Niu and Steve Phelps for their support and helpful discussions that have contributed to the findings of this paper.

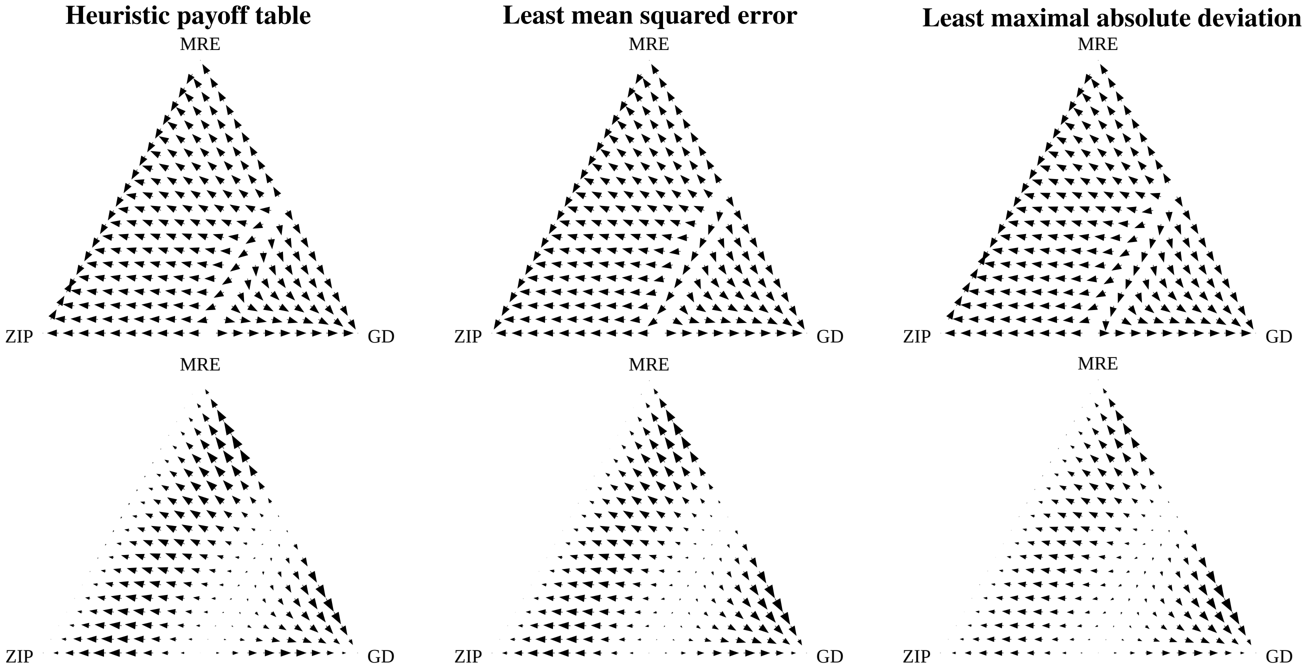

Figure 3. Comparison of the original replicator dynamics from the heuristic payoff table (left) to those from the normal form game approximations by least mean squared error (center) and minimized maximal absolute deviation (right) in the clearing house auction with 6 agents. The top row shows a directional field plot which allows to determine the basins of attraction by inspection. The bottom row shows a force field plot where the length of an arrow is proportional to the length of $\dot{x}$.

References

[1] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical report, Hewlett-Packard Research Laboratories, 1997.

[2] I. Erev and A. E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, 88(4):848–881, 1998.

[3] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29, January 1998.

[4] M. W. Hirsch, S. Smale, and R. Devaney. *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. Academic Press, 2004.

[5] T. B. Klos and G. J. van Ahee. Evolutionary dynamics for designing multi-period auctions (short paper). In *Proceedings 7th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*. IFAAMAS, 2008.

[6] J. McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–162, Summer 1994.

[7] J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and effciency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, 2001.

[8] S. Parsons, J. Rodriguez-Aguilar, and M. Klein. A bluffer's guide to auctions. Technical report, Center for Coordination Science, MIT, 2004.

[9] S. Phelps. Java auction simulator api. http://www.csc.liv.ac.uk/ sphelps/jasa/, 2005.

[10] S. Phelps, M. Marcinkiewicz, and S. Parsons. A novel method for automatic strategy acquisition in n-player non-zero-sum games. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 705–712, Hakodate, Japan, 2006. ACM.

[11] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[12] J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1993.

[13] M. A. Satterthwaite and S. R. Williams. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *The Review of Economic Studies*, 56(4):477–498, 1989.

[14] K. Tuyls and S. Parsons. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416, 2007.

[15] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In P. Gmytrasiewicz and S. Parsons, editors, *Proceedings of the Workshop on Game Theoretic and Decision Theoretic Agents*, pages 109–118, 2002.